DOCUMENT RESUME

ED 229 274 SE 041 559

AUTHOR	Reif, F.
TITLE	How Can Chemists Teach Problem Solving? Suggestions Derived from Studies of Cognitive Processes. Working Paper ES-17.
INSTITUTION	California Univ., Berkeley. Dept. of Physics.
SPONS AGENCY	National Science Foundation, Washington, D.C.
PUB DATE	82
GRANT	SED-79-20592
NOTE	33p.; Paper presented at the meeting of the American Chemical Society (Las Vegas, NV, March 1982). For related documents, see SE 041 560-561.
PUB TYPE	Reports - Research/Technical (143) -- Speeches/Conference Papers (150)
EDRS PRICE	MF01/PC02 Plus Postage.
DESCRIPTORS	*Chemistry; *Cognitive Processes; *College Science; Engineering; Higher Education; Knowledge Level; Mechanics (Physics); Physical Sciences; *Problem Solving; Recall (Psychology); Science Education; *Science Instruction; Student Behavior; *Teaching Methods
IDENTIFIERS	National Science Foundation; Science Education Research

ABSTRACT

Several central ideas emerging from a systematic approach to teaching problem-solving in the quantitative sciences (chemistry, physics, engineering) are discussed. Areas addressed include: differences between teaching and performance, between naturalistic and effective functioning, and between detailed observations and gross statistical data; insights derived from naturalistic studies, focusing on preexisting knowledge of students, tacit knowledge of experts, and significant differences between problem-solving behaviors of students and of experts; and kinds of procedures and knowledge essential for good human problem-solving performance, pointing out general issues addressed by any theoretical model of good problem-solving and discussing characteristics of the knowledge base containing knowledge about a specific domain. Problem-solving procedures considered include initial problem description, synthesis of the problem, and assessment/improvement of the solution. Current problem-solving activities in science teaching (focusing on student behaviors and instructional practices) are addressed, followed by a discussion of improved methods for teaching problem-solving. These methods include teaching explicitly and separately the various kinds of knowledge essential for good problem-solving performance (including knowledge of how to describe problem effectively), procedures useful for making judicious decisions in search for solutions, procedures for assessing solutions for correctness/optimality, and methods for organizing large amounts of knowledge so information can be easily recalled/remembered. (JN)

University of California, Berkeley, California 94720

Physics Department and Group in Science and Mathematics Education

Working Paper ES-17

HOW CAN CHEMISTS TEACH PROBLEM SOLVING?

SUGGESTIONS DERIVED FROM STUDIES OF COGNITIVE PROCESSES

F. Reif

This article is based on an invited paper presented at the meeting of the American Chemical Society, Las Vegas, March 1982. Some of the underlying work was partially supported by grant #SED 79-20592 from the National Science Foundation.

HOW CAN CHEMISTS TEACH PROBLEM SOLVING?

SUGGESTIONS DERIVED FROM STUDIES OF COGNITIVE PROCESSES*

F. Reif

Physics Department and Group in Science and Mathematics Education,

University of California, Berkeley, California 94720

Scientists do not behave scientifically in all domains. Thus we pursue our own discipline (e.g., chemistry) analytically and systematically, seek to develop a theoretical understanding of underlying processes, and try to achieve practical goals (e.g., chemical syntheses) on the basis of our theoretical insights. On the other hand, we commonly approach tasks outside our own discipline (even chemically related tasks such as gardening or nutrition) by the seats of our pants, content to rely on rules of thumb and on common-sense notions of questionable validity. Usually we tend to be equally unscientific in our teaching.

A serious interest in teaching scientific problem solving warrants, however, a more systematic approach. Not only is problem solving of crucial importance in any science if students are to achieve the ability to deal flexibly with diverse and novel situations. But problem solving, particularly in well-developed scientific disciplines, is also a highly complex intellectual task. Hence one cannot expect to achieve much success in teaching effective problem-solving skills unless one approaches the teaching enterprise from a systematic and scientific point of view. Such a scientific approach is not only required to achieve practical teaching effectiveness, but has also intrinsic intellectual interest as a field study in its own right. Indeed, recent years have seen substantial progress in our understanding of complex intellectual processes, as these have been studied in exciting new fields such as information-processing psychology or artificial intelligence [1-3].

My aim in this paper is to point out some central ideas emerging from a systematic approach to teaching scientific problem-solving skills. My comments will be generally applicable to problem solving in quantitative sciences, such as physics, chemistry, or engineering. (I myself have done most of my own work in the context of physics.) Within the limited time and space available to me, I shall focus my attention selectively on some major points and shall deliberately slight many important details.

Rudiments of a Systematic Approach

Teaching or learning can be viewed as a transformation process analogous to a chemical reaction of the form

$$S_i \to S_f$$

In this process a student S_i, in an initial state where he or she cannot do certain things (such as problem solving) is supposed to be transformed into a student S_f in a final state where he or she can do these things effectively.

An analysis of this transformation process reveals some basic issues which can usefully be distinguished and studied separately.

Teaching versus performance

It is clear that one must first have a good understanding of performance, both in the initial and final states, before one can systematically address tasks of learning or teaching. In particular, one must first understand how students, before any instruction, approach problems. Then one must understand how students, after instruction, are expected to perform, in order to achieve effective problem solving, e.g., what thought processes they are expected to use, how they are expected to organize their information, etc. (Indeed, it is a research challenge to understand theoretically such underlying cognitive processes leading to good human problem-solving performance.) Only after one has achieved a good understanding of initial performance, and of the desired final performance, can one hope to teach students how to become effective problem solvers.

Naturalistic versus effective functioning

Useful information about learning or teaching can obviously be obtained by observing and studying the performance of actual novice students and of actual experts. However, such naturalistic or "descriptive" studies have only limited interest. By contrast, a more general question, transcending a mere concern with naturalistic functioning, asks how effective functioning comes about. For instance, purely naturalistic studies of flying, by observing birds and insects, may lead to an understanding of how flying is achieved by flapping wings. However, asking the more general question about effective flying may lead to the realization that flying can be achieved even better with fixed wings, as modern airplanes do.

One can thus fruitfully ask the following general question, going beyond the bounds of descriptive naturalistic studies. What are underlying thought processes leading to good human performance (such as problem solving) without necessarily simulating what actual experts do? This more "prescriptive" question is both scientifically interesting and highly germane to practical instruction [4]. In particular, this question does not make the unwarranted assumption that actual experts always perform optimally. Furthermore, since it does not restrict inquiry to naturally occurring modes of functioning, it allows room for human invention and design. For instance, although models of good performance may be suggested by observing the behavior of actual experts, they may also be suggested by purely theoretical analyses.

The more general prescriptive question about good performance is also centrally important for teaching tasks or any attempts to improve human performance. In particular, it is a mistake to believe that good performance, to be achieved by students as a result of instruction, can merely mimic what experts actually do. Indeed, students must often be taught explicit processes to achieve performance which actual experts can do almost automatically because they immediately recognize situations familiar to them as a result of years of experience.

Detailed observations versus gross statistical data

To understand the underlying thought processes leading to good problem-solving performance, it is essential to observe in detail the thought processes of individual persons. By contrast, statistical data derived from test scores on many persons are of much less utility because they provide only very gross information. These comments are not

intended to denigrate the utility of statistical data. But, as somebody once said: "Statistics are like a bikini bathing suit. What they reveal is suggestive, but what they conceal is vital." There is a great deal of wisdom in this quotation.

Insights Derived from Naturalistic Studies

Before discussing in greater detail some of the thought processes needed for effective scientific problem solving, it is worth mentioning some insights derived from detailed observations of the naturally occurring problem-solving behavior of novice students and experts.

Detailed data of this kind can be obtained by asking individual persons to talk out loud about their thought processes while they are solving problems. The transcript of a person's tape-recorded verbal statements, together with the person's written work, constitutes then a "protocol" which provides a rich soucre of data about the person's thought processes. Needless to say, even such a detailed protocol reveals only a small part of a person's thoughts since many of these are not overtly verbalized. Nevertheless, such protocols provide data much more useful and detailed than would be obtained by test results, questionnaires, or other similar gross measures.

Let me then briefly mention some important results, derived from such detailed observations of novice students and experts [5-12] and point out some of their implications.

Preexisting knowledge of students

The observations indicate that novice students possess complex conceptual structures derived from prior experience and from informal cultural transmission. These conceptual structures are useful to explain and predict many of the phenomena encountered in daily life. However, unlike scientific conceptual structures, they are often ambiguous, vague, inconsistent, and not accurately predictive.

One implication of these observations is that the learning of a science involves much more than the acquisition of new knowledge by a blank mind. Instead, it involves a substantial restructuring of pre-existing knowledge, a restructuring where new knowledge must compete with a student's previous knowledge and familiar way of thinking. It is thus scarcely surprising that adequate restructuring can be a difficult and time-consuming process prone to many errors and confusions.

Another implication of these observations is that the modes of learning required in science (modes which require unambiguity, precision, and great care that all language is clearly related to observations) are quite unlike modes of learning familiar from daily life. Such new modes of learning are, therefore, quite difficult to acquire without carefully designed instruction.

Tacit knowledge of experts

Detailed observations indicate that experts possess knowledge which is remarkably large and well-organized. Much of this knowledge is "tacit", i.e., used automatically without any conscious awareness. Yet this tacit knowledge is essential to good performance and sometimes quite subtle.

One implication of these observations is that the explication of such tacit knowledge is an important and challenging task which can reveal much about the nature of expertise.

Another implication is that science teaching is often of limited effectiveness for the simple reason that much essential knowledge is never explicitly taught, because it is not even apparent to the teachers themselves.

Significant differences between experts and novices

Observations reveal that significant differences do indeed exist between the problem-solving behavior of novice students and of experts. For example, novice students usually try to assemble problem solutions by proceeding, in linear sequential fashion, to piece together various mathematical formulas. By contrast, experts often approach problems by using qualitative arguments and seemingly vague language, thus formulating plans which only later get refined into more mathematical language.

These observations show that experts' superior performance is not merely due to their large store of accumulated knowledge, but also to problem-solving strategies more effective than those used by students. As we shall see, such expert strategies are also theoretically expected to be more powerful and some of them should be teachable to students.

General Analysis of Effective Problem Solving

After these comments about information derived from naturalistic observations, let me turn to a more general analysis of the kinds of procedures and knowledge essential for good human problem-solving performance [13,14]. I shall begin by pointing out some general issues which must be addressed by any theoretical model of good problem solving. Then I shall examine some of these topics in slightly greater detail.

According to an analysis developed by myself and some coworkers, problem solving involves some general problem-solving procedures used in conjunction with a knowledge base containing particular knowledge about a specific domain (such as mechanics, or thermodynamics, or electric circuits, etc.). The general problem-solving procedures decompose the problem-solving process into several successive stages which address the following subproblems: (a) How can one initially describe and analyze a problem so as to facilitate the subsequent search for its solution? (b) How can one synthesize a solution of the problem, using appropriate planning and subsequent implementation, by making judiciously the many decisions needed to find a path to the solution? (c) How can one finally test the resulting solution to ascertain whether it is correct and reasonably optimal, so that suitable improvements can be made?

The preceding procedures are to be used in conjunction with a knowledge base containing specific knowledge about the particular domain of interest. Any such knowledge base must have general characteristics which facilitate the implementation of the preceding procedures, characteristics which must also be specified by a model of good performance. For example, what types of knowledge should be included in such a knowledge base? What kinds of ancillary knowledge must accompany any concepts or principle so that it becomes effectively usable and can thus serve as a functionally useful conceptual building block? Finally, how must the entire knowledge base be organized so that large amounts of information can be easily remembered and appropriately retrieved in complex problem-solving contexts?

An understanding of how good human problem-solving performance can be achieved requires answers to <u>all</u> the preceding questions. In the

following paragraphs I shall merely outline some of the major ideas which have emerged in our work addressing these questions.

Problem-solving Procedures

Initial problem description

The manner in which a problem is initially described is crucially important since it can determine whether the subsequent solution of the problem is easy or difficult--or even impossible. The crucial role of the initial description of a problem is, however, easily overlooked because it is a preliminary step which experts usually do rapidly and automatically without much conscious awareness.

A model of effective problem solving must thus, in particular, specify explicitly procedures for generating a useful initial description of any problem. The first stage of such a description procedure aims to generate a "basic description" of a problem. This is achieved by using general domain-independent knowledge to put the problem into a form where it is readily understandable to the problem solver. Thus the basic description summarizes the information specified and to be found, introduces useful symbols, and expresses available information in various useful symbolic forms (e.g., in verbal statements as well as in diagrams). Figure 1 illustrates an example.

Insert Figure 1 about here

The next stage of the description procedure is more complex. It aims to generate a "theoretical description" of the problem, i.e., a description which deliberately aims to redescribe the problem in terms

ORIGINAL PROBLEM STATEMENT

A MAN SITS AT THE TOP OF A SMOOTH HEMISPHERICAL DOME, OF RADIUS R, COVERING A FACTORY ON HORIZONTAL GROUND. IF THE MAN STARTS SLIDING, AT WHAT HEIGHT ABOVE THE GROUND DOES HE SLIDE OFF THE DOME?

BASIC DESCRIPTION

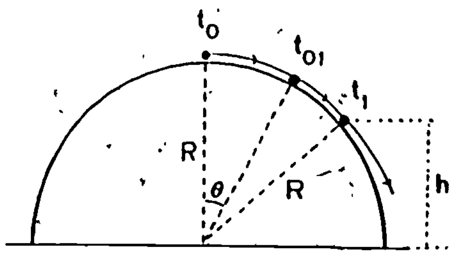

time t_0 : man at top

time t_1 : man slides off

(time t_{01} : any time between)

GOAL: $h = ?$

Fig. 1: Original statement and basic description of a mechanics problem.

of the special concepts provided by the knowledge base for the relevant knowledge domain. The resulting problem description greatly facilitates the subsequent search for a problem solution since all principles in the knowledge base are expressed in terms of these special concepts and become thus readily accessible.

Hence the knowledge base for any particular domain is especially useful if it includes explicit rules specifying how to describe any situation encountered in this domain. These rules should specify how to identify, in any problem situation, those entitities of prime interest in this domain, what special concepts should be used to describe these entities, what properties of these special concepts can be exploited, and how to check that the resulting description is consistent with known principles. [15]

For example, the knowledge base for the science of classical mechanics can usefully be accompanied by explicit rules specifying explicitly how to describe any problem in mechanics. These rules specify that the entities of interest are particles (or systems consisting of several such particles). They specify that the <u>motion</u> of any such particle should be described by special concepts such as "position", "velocity", and "acceleration". They also specify that the <u>interaction</u> between such particles should be specified by special concepts such as "forces" or "potential energies". The description rules also specify how to exploit the properties of these concepts. For instance, they specify that forces can be systematically enumerated by first considering long-range forces (such as gravity) and then identifying the short-range forces on any particle (by noting all objects which touch the given particle). Furthermore, they specify how these various kinds of forces (such as gravitational forces by the earth, forces by strings, etc.) depend on the characteristics and positions of the interacting particles. Finally, they specify that a problem description must be consistent with motion principles, e.g., that the acceleration of any particle must have the same direction as the total force on it. Figure 2 illustrates such a theoretical description of the problem previously described in Figure 1.

Insert Figure 2 about here

Such description rules are considerably more explicit than those usually taught or found in textbooks. But in experiments, performed by Heller and myself, we showed convincingly that students, when induced to use such description rules, avoid almost all common errors (e.g., omitting relevant forces or introducing extraneous forces due to non-existent objects) and generate problem descriptions leading to successful solutions [15].

Similar description rules can be explicitly formulated for other knowledge domains. For example, in thermodynamics the entities of interest are macroscopic systems consisting of very many atomic particles. Such isolated systems are to be described by special concepts, such as "internal energy" and "entropy", and this description can exploit the special properties of these concepts. Once again, the mere redescription of thermodynamic problems in terms of these particlar concepts greatly facilitates their subsequent solutions.

Synthesis of a solution.

Once a problem has been described, one can turn to the task of

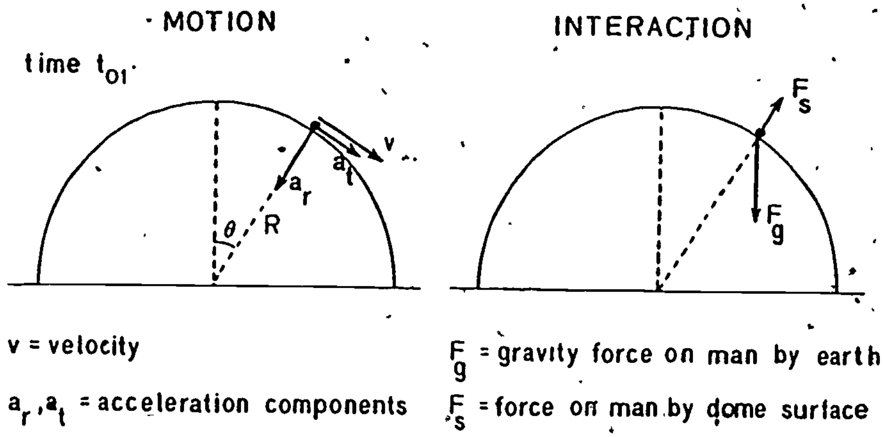

Fig. 2: Theoretical description of the problem of Figure 1.

constructing its solution. This task is difficult because the search for a solution requires decisions among many possible alternatives, only a few of which lead to the desired goal. Hence a model of effective problem solving must specify how to make judicious decisions to find an efficient path to the desired goal.

Merely explicating the alternatives to be considered in making a decision can already be useful, even without specifying how a choice between these alternatives is to be made. For example, a typical alternative is of the form. "What principle should be applied to what system at what time (or between what times) with what description?" An explicit awareness that this is the kind of decision to be made helps to identify a limited number of component alternatives worthy of consideration and thus simplifies appreciably the decision process. For example, the specified form of such a decision may help to focus a student's diffuse thinking on "applying the first law of thermodynamics, to a system consisting of a particular cylinder and enclosed gas, between the states A and B, described in terms of pressure and volume".

Much could be said about how to make judicious decisions among promising alternatives once they have been identified. But instead of discussing several useful decision methods in greater detail, let me merely make a few comments about a powerful general strategy used to facilitate problem solving, the strategy of progressive refinements.

This strategy may usefully be illustrated by an analogy, the problem of painting a picture. One painting strategy would be to paint successively, in complete detail, every adjacent square inch of the picture until the total picture is completed. The other strategy consists of first making a rough sketch of the entire picture, then elaborating this

sketch by adding more detailed lines, then elaborating further by adding more detailed color information, etc. The second strategy, which proceeds by progressive refinements, is far more useful because it allows one to make crucial major decisions first without burdensome or distracting details. These major decisions can then be used as guides to make further decisions at a more detailed level, and so on until all details have been worked out.

Similarly, in scientific problem solving it is useful to make first a few major decisions by describing only the major features of a problem in gross qualitative terms (using vague language or pictures). Such a solution plan, outlined at a gross level of description, may then be used as a guide to construct a solution in more precise and mathematical language. Indeed, as mentioned previously, experts (unlike novice students) commonly use such methods of progressive refinement with great effectiveness.

As a simple illustration of the preceding remarks, consider the situation illustrated in Figure 3 which shows a current flowing through two joined wires of different diameters. The problem is to find the potential at the junction point if the potentials at the ends of the wires are specified. Experts commonly approach this problem by stating verbally that the potential drop in each wire depends on its resistance, and that the resistance of each wire depends on its geometric properties. Only after such qualitative remarks, which basically constitute a rough plan for a solution, do the experts begin to specify in greater detail just how the resistance of each wire depends on particular geometric properties, such as length and cross-sectional area. This strategy of progressively refined descriptions stands in sharp contrast to what novice students usually do, namely trying to construct a solution by merely writing down various equations at a mathematically detailed level of description. (Indeed, many students get lost in a morass of several such equations in several unknowns, and thus find the preceding simple problem difficult.)

Insert Figure 3 about here

Assessment and improvement of solution

Once a problem solution has been obtained, it is important to assess whether it is correct and reasonably optimal, so that suitable improvements can be made. Various tests, designed to assess whether a solution is correct, can readily be formulated in explicit form. For example, one of the most useful of such tests is a consideration of simple special cases (particularly of extreme cases) which must be consistent with a general solution of a problem. Other useful tests can also be easily specified. Indeed, many such tests are quite straightforward and familiar to experts, although they are often not explicitly taught to students.

Characteristics of the Knowledge Base

As mentioned previously, the knowledge base for a particular domain must have general characteristics which facilitate the preceding procedures used to describe problems and search for their solutions. Let me discuss slightly more fully some of these important characteristics.

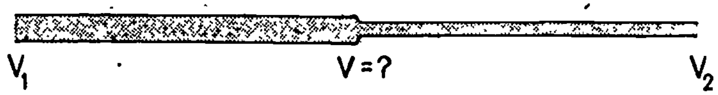

Fig. 3: Current flowing in joined wires of different diameters.

Functional conceptual building blocks.

Particular special concepts and principles are, from a strictly logical point of view, the essential building blocks of the knowledge used to make scientific predictions and to solve problems. However, the mere definition of a concept, or mere statement of a principle, is psychologically and practically almost worthless unless it is accompanied by ancillary knowledge needed to make the concept or principle effectively usable. Only then does the concept or principle become a functionally useful conceptual building block suitable for the synthesis of more complex problems.

An analysis, presented elsewhere [16], shows that the ancillary knowledge needed to make any concept or principle effectively usable is remarkably large and sometimes subtle. Thus the knowledge needed to interpret a concept or principle includes that required to specify the concept or principle by descriptive statements and by detailed procedures needed to identify the concept; it includes the knowledge needed to apply the concept in various specific kinds of instances; and it includes explicit warnings about likely errors and the requisite knowledge to discriminate them from correct situations. The ancillary knowledge includes also familiarity with some basic implications and applications of a concept or principle, as well as explicit guidelines specifying the conditions when the concept is likely to be useful.

Such ancillary knowledge, required to make a concept or principle effectively usable, is routinely possessed by any expert. However, much of this knowledge is tacit and often not explicitly taught. Indeed, deficiencies in such ancillary knowledge lead to many of the common misconceptions and errors committed by students.

The preceding comments indicate that it is important to make explicit, and also to teach deliberately, the ancillary knowledge required to make any concept or principle effectively usable. In this way one can ensure the possession of concepts and principles which are functionally useful and which thus provide a necessary (although not sufficient) prerequisite for effective problem solving.

Knowledge Organization

The organization of available knowledge is of crucial importance, particularly if the knowledge is large. Indeed, only if information is effectively organized, can it be easily remembered and appropriately retrieved in complex contexts. (For example, although every folder in a file cabinet may contain valuable information, an unorganized collection of such folders would make the information available in the file cabinet nearly inaccessible and thus almost useless.) A theoretical model of effective problem solving must thus specify explicitly the manner in which concepts, principles, and rules should be effectively organized so as to facilitate their ready retrieval and use.

A particular form of knowledge organization, highly useful in many cases, is one which is hierarchically structured at successive levels containing increasingly more detailed information [17]. This form of organization may be illustrated by a familiar analogy, the manner in which geographical information is organized into maps of different scales. Thus geographical information about the entire United States can be summarized in a map showing very little detailed information and only gross features. Then certain of these gross features (e.g., certain geographical regions) can be elaborated more fully in other maps (e.g., maps of particular states) showing more details. Such maps can, in turn, be elaborated more fully in other maps which show still more details. In this way, proceeding by successive elaboration, it is possible to accommodate any amount of detail in a fashion which does not obscure the major features and which uses these major features as aids to gain access to more details.

Scientific knowledge about any domain, such as mechanics or thermodynamics, can very effectively be organized in similar hierarchical ways which facilitate the remembering and retrieval of all the relevant knowledge. It would take me too far afield to show here such a hierarchical organization of an entire scientific domain. Let me therefore merely illustrate the utility of such a hierarchical organization in a very simple, but common, case.

Consider an argument (e.g., the derivation of some scientific result) which starts from certain premises to arrive at certain conclusions. Such an argument can be organized in various alternative ways. A purely linear organization of this argument might consist of a dozen or so successive steps, each well explained and following logically from the preceding one. On the other hand, a hierarchical organization of the same argument might consist of four major steps, grossly described, which summarize the entire argument; each of these may then, in turn, be elaborated into several more detailed steps, as indicated schematically in Figure 4.

Insert Figure 4 about here

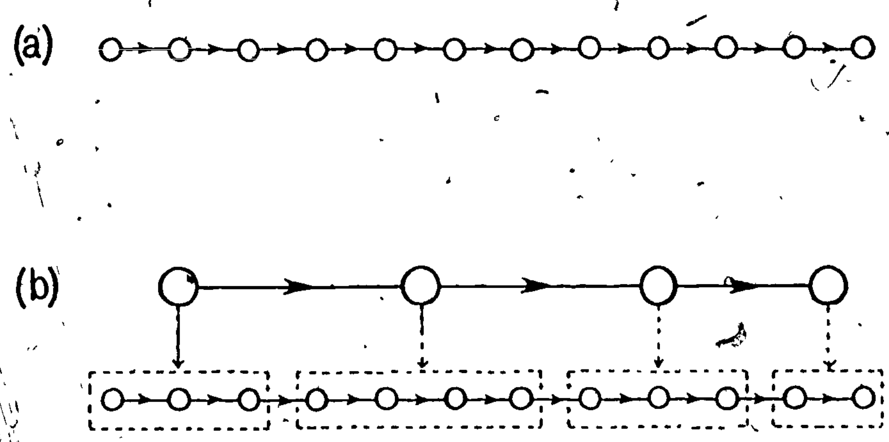

Fig. 4: Schematic illustration of steps in an argument. (a) Linear organization of the argument. (b) Hierarchical organization of the same argument.

Although the hierarchical organization of the argument contains the same information, it is expected to be much more useful. Indeed, experiments done by Eylon and myself have shown that an argument, presented and learned in such a hierarchical form, can be more easily remembered, more easily modified if some premises are changed, and more easily corrected if errors are made [17].

Experimental Tests of Theoretical Ideas

In the preceding sections I have discussed various theoretical ideas specifying procedures and forms of knowledge leading to good human problem solving. As mentioned previously, such theoretical ideas need not necessarily simulate the actual behavior of experts. The central question is rather whether human beings, acting in accordance with such theoretical ideas, do indeed achieve good performance. To address this question, my collaborators and I have used the following experimental paradigm to test models of good problem-solving performance: (1) We use carefully controlled experimental conditions to induce individual persons to perform in accordance with a specified model of good performance (e.g., we induce persons to follow step-by-step directions specifying how to generate useful initial descriptions of problems). (2) Then we observe in detail the resulting performance of these persons and assess whether this performance exhibits the predicted characteristics and is effective.

Such experiments can indeed be successfully implemented in practice and can yield very valuable information, both to test theoretical models and to suggest improvements in them. For example, some of the previously mentioned theoretical ideas about effective problem description and

about effective knowledge organization were tested by us in experiments of this kind [15, 17].

Current Actualities in Science Teaching

The analysis on the preceding pages outlines some useful procedures and kinds of knowledge facilitating scientific problem solving. How do these insights about effective problem solving relate to the state of affairs prevalent in the actual teaching of science? In particular, what do students commonly do when approaching problems? And what do instructors commonly do when trying to teach scientific problem solving?

Common student behaviors

Informal observations, as well as some systematic studies [11, 12] indicate that typical students in basic college-level science courses tend to approach problems in a manner that differs appreciably from the preceding precepts for good problem-solving. The following are some examples.

Students, faced with a problem, try to tackle immediately the task of constructing a solution. They spend rather little time beforehand to generate a careful description of the problem or to plan a solution. Furthermore, they spend little or no time afterwards to assess whether the solution found by them is correct or as simple as it might be; nor do they try to extract from a solution useful knowledge that might help them to deal with future problems.

Students usually try to find the solution of a science problem by assembling various facts and formulas in a linear sequential fashion, rather than by using methods of progressive refinement. Students tend to behave similarly in other problem-solving domains. For example, when trying to write English prose, novice students tend to proceed by generating successive sentences and paragraphs, rather than by constructing progressively refined outlines and drafts. Similarly, when trying to write computer programs, unsophisticated students tend to proceed by writing successive lines of code, rather than by progressively elaborating flow charts or high-level procedural specifications. It is not surprising that students behave in this way. Faced with a task which ultimately consists of a sequence of steps, the most primitive strategy is to generate the steps one-by-one in sequential order. By contrast, a strategy of progressive refinements, which first generates more abstract steps to be ultimately elaborated into the steps of actual interest, is a more indirect method and its superiority is not apparent to unsophisticated students.

Students tend to place a great emphasis on remembering and using various facts and mathematical formulas, without trying to embed these in a rich framework of qualitative knowledge. Accordingly, students may be able to answer some quantitative questions by merely manipulating some mathematical formulas, but may be quite unable to answer simple qualitative questions of a similar kind. Furthermore, students seldom use qualitative knowledge to plan solutions or to check whether results obtained by them make any sense.

In trying to learn problem solving, students pay much more attention to the product than to the process. Thus students are mostly interested in the answers to problems or in worked-out solutions of them. But they often fail to realize that the most important aspect of problem solving

involves the decision processes which lead to good solutions, decision processes which prevent one from getting stuck or going off in wrong directions.

Finally, students' acquired knowledge about a scientific domain is often poorly structured, consisting of little more than loosely organized collections of miscellaneous facts or lists of such facts. Indeed, students rarely appreciate the great importance of organizing their knowledge with great care so as to make it effectively usable. Nor are they familiar with forms of organization, such as hierarchical organizations, potentially useful for making their knowledge more coherent, more easily remembered, and more readily retrievable.

In summary, it is clear that students' problem-solving skills are rather primitive and leave ample room for improvement through specific teaching.

Common teaching practices

Common teaching methods, used by instructors or textbooks to teach scientific problem-solving skills, rely predominantly on presenting information, showing prototypical examples of worked-out problems, and providing students with practice in solving similar kinds of problems. In many ways such teaching methods are more primitive than those used in simpler domains, such as sports or musical performance. Indeed, in those domains teachers try to analyze in detail important components needed for good performance (e.g., how to hold a violin, how to move the wrist when bowing, etc.) Then they strive to teach explicitly these important components and to integrate them into good overall performance. These teaching methods are thus more systematic and analytic than those based on mere exemplification and practice.

Not only do the teaching methods used in scientific instruction seem fairly primitive, there is also evidence that they are often rather inefficient and ineffective. For example, several recent studies [5-10] show that many students, after having successfully completed college-level science courses, may nevertheless exhibit gross errors or misconceptions and be quite deficient in simple problem-solving skills.

Common teaching practices are not only of limited effectiveness, but may even be dysfunctional or deleterious. For example, quite often instructors, in their presentations and examinations, emphasize factual knowledge considerably more than reasoning processes. Furthermore, teachers and textbooks, in a narrowly conceived pursuit of scientific precision, frequently tend to stress formal knowledge and mathematical descriptions. Correspondingly they neglect, or even discourage, qualitative understanding and modes of qualitative reasoning which (as pointed out previously) are very powerful aids for planning problem solutions and are commonly used by actual experts. As another example, scientific arguments and prototypical problem solutions are usually presented as linear sequences of detailed steps, rather than in more hierarchical forms which (as pointed out previously) are much more useful for remembering and generalizing information.

The preceding comments reveal that common teaching practices do not reflect much insight into effective problem-solving processes and leave much room for improvement.

Improved Methods for Teaching Problem Solving

The discussion in the preceding pages suggests that the teaching of scientific problem solving could be substantially improved if the task were approached more scientifically and systematically. In particular, such teaching should be based upon an adequate understanding of how good problem solving is achieved, i.e., upon an analysis of the kind sketched in the preceding pages [13, 14].

Such an analysis suggests that one teach explicitly and separately the various kinds of knowledge essential for good problem-solving performance, e.g., knowledge of how to describe problems effectively, procedures useful for making judicious decisions in the search for solutions, procedures for assessing solutions for correctness and optimality, methods for organizing large amounts of knowledge so that information can be easily remembered and retrieved, etc. Appropriate teaching methods must then be used to integrate these important components so that students can use them jointly in coherent fashion. Finally, one must ensure that students learn to use these methods habitually and automatically. Clearly, adequate practice is needed to achieve these ends, but the <u>right kinds</u> of practice based on an explicit understanding of underlying mechanisms of good problem solving.

The preceding general suggestions can be elaborated into practical teaching programs. For example, I myself have tried to do so in a special course on physics problem solving. I have also explored a preferable course format where students are explicitly taught scientific conceptual and problem-solving skills in a "workshop" accompanying the first quarter of an introductory college-level physics course. Finally, St. John and I have even attempted to incorporate a few of these ideas in laboratory instruction [18]. Such practical teaching efforts show definite promise, but need appreciable care and thought for proper implementation. In particular, they require painstaking attention to many important details which I have not been able to mention in the brief survey presented in the preceding pages. Such teaching efforts could also benefit from the exploitation of more individualized forms of instruction implemented by programmed teaching materials or computer-aided instruction.

Finally, one should not expect miracles, even if the teaching of problem solving is approached from a systematic and scientific point of view. After all, such teaching efforts deal with very complex cognitive skills. But, by approaching such teaching tasks systematically, one may hope to achieve the advantage inherent in scientific approaches in other fields, namely cumulatively increasing knowledge and understanding, opportunities to learn from successes as well as from failures, and increased practical effectiveness based upon validated theoretical insights.

FOOTNOTE

* This article is based on an invited paper presented at the meeting of the American Chemical Society, Las Vegas, March 1982. Some of the underlying work was partially supported by grant #SED 79-20592 from the National Science Foundation.

References

(1) Anderson, J. R., "Cognitive Psychology and its Implications", W. H. Freeman, San Francisco, 1980.

(2) Newell, A., and Simon, H. A., "Human Problem Solving", Prentice-Hall, Englewood Cliffs, N.J., 1972.

(3) Winston, P. H., "Artificial Intelligence", Addison-Wesley, Reading, Mass., 1977.

(4) Reif, F., in "Problem Solving and Education: Issues in Teaching and Research" edited by Tuma D., and Reif F., (Lawrence Erlbaum Associates, Hillsdale, N.J., 1980).

(5) Lochhead, J. and Clement J., (eds.), "Cognitive Process Instruction", Franklin Institute Press, Philadelphia, 1979.

(6) Viennot, L., European J. of Science Ed., 1, 205 (1979).

(7) Trowbridge, D. E., and McDermott, L.E., Am. J. Phys. 48, 1020 (1980) and Am. J. Phys. 49, 242 (1981).

(8) McCloskey, M., Caramazza, A., and Green, B., Science, 210, 1139 (1980).

(9) Champagne, A. B., Klopfer, L. E. and Anderson, J. H., Am. J. Phys. 48, 1071 (1980).

(10) Clement, J., Am. J. Phys. 50, 66 (1982).

(11) Larkin, J. H. and Reif, F., European J. of Science Education, 1, 191 (1979).

(12) Larkin, J., McDermott, J., Simon, D. P., and Simon, H. A., Science 208, 1335 (1980).

(13) Reif, F. and Heller, J. I., "Knowledge Structure and Problem Solving in Physics" (Report ES-12, Physics Department, University of California, Berkeley, 1981. To be published in the Educational Psychologist.)

(14) Reif, F., The Physics Teacher, 19, 310 (1981).

(15) Reif, F. and Heller, J. I., "Cognitive Mechanisms Facilitating Human Problem Solving in Physics." (Report ES-14, Physics Department, University of California, Berkeley, 1982; presented at annual meeting of American Educational Research Association.)

(16) Reif, F., "Making Scientific Concepts and Principles Effectively Usable: Requisite Knowledge and Teaching Implications." (Report ES-13, Physics Department, University of California, Berkeley, 1982.)

(17) Eylon, B. and Reif, F., "Effects of Knowledge Organization on Task Performance" (Unpublished Report, Physics Department, University of California, Berkeley, 1979; presented at annual meeting of American Educational Research Association.)

(18) Reif, F. and St. John, M., Am. J. Phys. 47, 950 (1979).